



The Machine Learning Pipeline on AWS

Course Overview

This course explores how to use the iterative machine learning (ML) process pipeline to solve a real business problem in a project-based learning environment. Students will learn about each phase of the process pipeline from instructor presentations and demonstrations and then apply that knowledge to complete a project solving one of three business problems: fraud detection, recommendation engines, or flight delays. By the end of the course, students will have successfully built, trained, evaluated, tuned, and deployed an ML model using Amazon SageMaker that solves their selected business problem. Learners with little to no machine learning experience or knowledge will benefit from this course. Basic knowledge of Statistics will be helpful.

Prerequisites

- Basic knowledge of Python programming language
- Basic understanding of AWS Cloud infrastructure (Amazon S3 and Amazon CloudWatch)
- Basic experience working in a Jupyter notebook environment

Target Audience

Developers

Solutions Architects

Data Engineers

Anyone with little to no experience with ML and wants to learn about the ML pipeline using Amazon SageMaker

Course Objectives

- Select and justify the appropriate ML approach for a given business problem
- Use the ML pipeline to solve a specific business problem
- Train, evaluate, deploy, and tune an ML model using Amazon SageMaker
- Describe some of the best practices for designing scalable, cost-optimized, and secure ML pipelines in AWS
- Apply machine learning to a real-life business problem after the course is complete

Duration

4 Days

Certifications

Certified Machine Learning
- Specialty

Contact Us

800.674.3550

2151 W. Hillsboro Blvd.
Suite 210
Deerfield Beach, FL 33442

Connect with us



Sign Up Today!





The Machine Learning Pipeline on AWS



Course Outline

Module 1: Introduction to data lakes

- Describe the value of data lakes
- Compare data lakes and data warehouses
- Describe the components of a data lake
- Recognize common architectures built on data lakes

Module 0: Introduction

- Pre-assessment

Module 1: Introduction to Machine Learning and the ML Pipeline

- Overview of machine learning, including use cases, types of machine learning, and key concepts
- Overview of the ML pipeline
- Introduction to course projects and approach

Module 2: Introduction to Amazon SageMaker

- Introduction to Amazon SageMaker
- Demo: Amazon SageMaker and Jupyter notebooks
- Hands-on: Amazon SageMaker and Jupyter notebooks

Module 3: Problem Formulation

- Overview of problem formulation and deciding if ML is the right solution
- Converting a business problem into an ML problem
- Demo: Amazon SageMaker Ground Truth
- Hands-on: Amazon SageMaker Ground Truth
- Practice problem formulation
- Formulate problems for projects

Checkpoint 1 and Answer Review

Module 4: Preprocessing

- Overview of data collection and integration, and techniques for data preprocessing and visualization
- Practice preprocessing
- Preprocess project data
- Class discussion about projects

Checkpoint 2 and Answer Review

Module 5: Model Training

- Choosing the right algorithm
- Formatting and splitting your data for training
- Loss functions and gradient descent for improving your model
- Demo: Create a training job in Amazon SageMaker

Module 6: Model Evaluation

- How to evaluate classification models
- How to evaluate regression models
- Practice model training and evaluation
- Train and evaluate project models
- Initial project presentations

Checkpoint 3 and Answer Review

Module 7: Feature Engineering and Model Tuning

- Feature extraction, selection, creation, and transformation
- Hyperparameter tuning
- Demo: SageMaker hyperparameter optimization
- Practice feature engineering and model tuning
- Apply feature engineering and model tuning to projects
- Final project presentations

Module 8: Deployment

- How to deploy, inference, and monitor your model on Amazon SageMaker
- Deploying ML at the edge
- Demo: Creating an Amazon SageMaker endpoint
- Post-assessment
- Course wrap-up

